Architecting Technology Transition Pathways: Insights from the Military Tactical Network Upgrade[1]

Zoe Szajnfarber,[2] Linda McCabe[3] and Amanda Rohrbach[4]

Abstract

As engineering systems become more and more complex, technology transition increasingly involves deploying an upgraded subsystem across a legacy network. This mode of upgrade presents new challenges for systems architects concerned with maintaining value over multiple infused technical changes. This paper explores the dynamics of technology transition in path-dependent infrastructure systems. It uses a model-based case study of the envisioned military Airborne Tactical Network (ATN) upgrade as a basis for developing guidelines for effective transition path design. Based on the natural diffusion dynamics of the system we identified an inherent tradeoff between upgrade cycle and sustained capability levels. In other words, assuming even weakly exponential growth in demand, there is a relationship between timing of infusion and longevity of benefit. As a result, a less capable upgrade, deployed expediently can do more good than a more sophisticated upgrade that can only be integrated in the next block upgrade. In addition, by conceptualizing the transition “path” as a design lever, two dimensions of problem decomposition can be exploited to mitigate transition barriers: (1) Self-contained sub-networks can provide a proving ground for full-system future benefits in order to mitigate stakeholder resistance; and (2) The technical system can be designed for evolvability, making it possible to stage deployment in the technical dimension as well.

1. Introduction

Today’s engineering systems are increasingly complex and interdependent. As part of this trend, the luxury of *green-field* design has become a rarity in many contexts. More often, systems are evolved piece-by-piece, rather than being replaced all at once. Upgrading/inserting new technology into a fielded complex system can be not only very expensive, but also take many years to complete. This reality has created a new, arguably more difficult, challenge for systems planners and architects. Where in the past, the challenge was to identify the set of capabilities that could best achieve the stakeholders’ objectives; today, many key design variables are pre-determined by the legacy system, and therefore outside of the systems architect’s control. Constrained by fixed elements of legacy architecture, systems engineers must now plan a *path for technology transition* that evolves the system to a desired state, with minimal disruption to ongoing operations. In this context, understanding the dynamics of the system, key stakeholder equities, and how

[1] Distribution A: Public Release This work was sponsored by the Assistant Secretary of Defense – Research and Engineering (ASD R&E) under Air Force Contract # FA8721-05-C-0002. Opinions, interpretations, recommendations and conclusions are those of the authors and are not necessarily endorsed by the United States Government.
[2] Assistant Professor of Engineering Management and Systems Engineering, George Washington University, zszajnfa@gwu.edu.Mazzuchi
[3] Technical Staff, Airborne Networks MIT Lincoln Laboratory, lmccabe@ll.mit.edu.
[4] Research Assistant, Technology and Policy Program, MIT, rohrbach@alum.mit.edu.

costs and benefits of differing transition strategies ripple across the system, is critical to selecting a transition strategy that will ultimately yield benefits to stakeholders as early in the transition as possible while controlling costs.

One modern example of this formulation of the so-called path dependent systems architecting problem is the intention of the U.S. Department of Defense (DoD) to upgrade its tactical communications network. Imminent capacity constraints have been consistently forecasted and there is general agreement on the desired future state: "It is DoD policy that…. Communications waveform development or modification, and the associated network, will implement Internet Protocol (IP) capability to the extent possible to enable net-centric interoperability" [Grimes, 2008]. However, the *path through which technology transition should be structured* is the subject of considerable debate.

Determining an appropriate transition pathway (i.e., the order in which upgrades should be deployed and the matching of upgrades to existing airframes) is complicated by multiple factors. Even just considering the airborne tactical network (ATN), upgrading the communications system means a software and/or hardware change to each of the 6000+ airframes currently in the inventory. These airframes are deployed all over the world and the process takes time. The specifics of how the upgrade can be implemented depend on characteristics of both the aircraft (i.e., specific size, weight and power issues, how near it is to retirement, the envelope available in which to make changes) and the technical radio system being infused. Some upgrades can be done in the field, while others require a return to a U.S.-based depot, or may not be feasible until the next block-upgrade leaves the assembly line.

Order matters because of the inherent lag between when costs are committed and benefits begin to accrue. Since only pairs of upgraded aircraft can use the improved channel, legacy technology will impose a disproportionate drag on the capability of the system. As a result, unlike in the commercial world, where there is a *first mover* advantage, early adopters of IP-based radios may wait years before they see the fruits of their investment. This is because costs of the upgrade must be born upfront – likely by individual services, and within specific programs – but benefits may take years to accrue. And, the lag between cost commitment and benefit accrual will be a non-linear result of how the transition is managed.

These types of transition problems are not expected to lessen in the future. For example, the FAA is currently struggling with encouraging new satellite-based technology adoption to support NextGen, the DOE is exploring policies to promote deployment of cleaner technologies in e.g., the coal industry and the cellphone sector is continuously upgrading to more modern network standards. In all these cases the system is characterized by legacy infrastructure being upgraded by changing out core components over time. The transitions have all proved more challenging than anticipated. The goal of this work is to provide better guidance on how to architect transition pathways, by staging the deployment in either technical or stakeholder dimensions. To that end, this paper uses a model-based case study of possible transitions of the ATN to develop design principles that can guide future transition planning.

The paper makes two specific contributions. First, the work illustrates the value of simple models in revealing change patterns in fielded systems. The model provides a basis for evaluating possible ATN upgrade paths and yields practical insights for decision makers in that space.   Second, by introducing the path-dependent systems upgrade problem – a previously under-studied but increasingly important class of system change – the paper sheds light on an important area of future systems engineering research. This initial work specifically demonstrates the importance of characterizing the inherent upgrade cycles of the system, and planning accordingly. Further, it shows that by conceptualizing the transition "path" as a design lever, two dimensions of problem decomposition can be exploited to mitigate transition barriers: (1) Self-contained sub-networks can provide a proving ground for full-system future benefits in order to mitigate stakeholder resistance; and (2) The technical system can be designed for evolvability, making it possible to stage deployment in the technical dimension as well.

The body of the paper is organized in four sections: Section 2 grounds the problem in related literature; Section 3 describes the system modeling approach and associated assumptions; Section 4 describes case-specific results; and Section 5 generalizes those results to other similar transition contexts.

2. Theoretical Basis for Architecting Technology Transitions

The literature explicitly on strategies for architecting technology transition in path dependent systems is fairy limited. Nonetheless there is substantial related literature describing sources of resistance to change, models of change at the policy level and evolvability/change in technical systems. This section synthesizes those insights as a basis for defining alternative transition strategies, which will be compared and evaluated in the rest of the paper.

For our purposes, transition pathways are defined as the set of steps taken to implement a change. For example, in the context of ATN, the relevant variables to order are which aircraft should receive the upgrade first (e.g., Navy vs. Air Force, or US-based fleet vs. those deployed in Asia) and what level of capability should be deployed (e.g., swap-out the whole communication system vs. start with a software patch). By path-dependent infrastructure, we consider systems that are not replaced as a unit. As a result, new infused modules/components must work with an existing infrastructure that is not also being replaced.

2.1 General Barriers to Change

One theme that runs across all political science and management scholarship on change processes is that bureaucratic systems inherently resist change [c.f., Rosen, 1994; Sapolsky, 1972]. The particular sources of resistance are somewhat system-dependent, but can broadly be characterized as kinds of "lock-in." The idea is that established firms/actors, who have the most relative power in the current system, also have the most to lose if the status quo changes. They will therefore use their power to resist change either actively [e.g., through regulatory lock-in [Stigler, 1971] or passively by pushing the incumbent approach well beyond its usefulness.

In commercial industry, this resistance has been explained in the technology management literature through several mechanisms. First, the attributes of a firm that enable it to optimize performance contradict the agility required to explore multiple novel approaches [Utterback and Abernathy, 1975; Utterback, 1994; Schumpeter, 1934]. Second, incremental changes tend to enhance the competences of established firms, while radical, or even architectural [Henderson and Clark, 1990] innovations tend to destroy those competences [Anderson and Tushman, 1990]. As a result, it is both against a firm's short-term self-interest to destroy its own competence, and also more difficult for that firm to see competence destroying opportunities [Henderson and Clark, 1990]. History is replete with examples of powerful incumbents failing to traverse technical discontinuities, and being replaced with entrepreneurial disruptors [Schumpeter, 1934; Christensen, 2003].

In the government sector, where agencies and services have nearly assured longevity, the dynamics can be somewhat different. The resistance to change is institutionalized through promotion pathways that reinforce existing preference [Rosen 1994]. In addition, change in the government setting often requires multiple autonomous entities to work together, even when interests are not completely aligned. For example, several studies of transition in the air transportation system [Mozdzanowska et al., 2008; Mozdzanowska 2008; Marais and Weigel, 2006] have highlighted the multi-stakeholder problem. The transition from radar- to GPS-based situational awareness will eventually benefit everyone, but the new operating procedures that enable time and fuel savings as well as safer operations can't take effect until all or most of the airframes have the new equipment. Since the cost-burden will be unequally born by airlines and private operators upfront, and significant changes in tasks will affect air traffic controllers first, achieving stakeholder consensus has proved a significant barrier to adoption.

Marais and Weigel (2006) characterize these barriers in terms of misalignment in the distribution of costs and benefit across stakeholders and most importantly, the lag between cost commitment and benefit accrual on the part of particular stakeholders. This framework is a valuable lens through which to evaluate resistance to change in the ATN adoption context.

2.2 Models of Change

Given the inherent barriers, change typically happens in one of two ways: through *incrementalism* or as a *shock* in response to a crisis. Incrementalism assumes that only small changes at the margins are ever possible, so most outcomes are the result of multiple incremental suboptimal changes [Lindblom, 1992]. A slightly less pessimistic view of this type of process suggests that a baseline level of muddling through is required to ensure that a particular solution is ready if an opportunity to act (e.g., a policy window) ever opens [Kingdon, 1984]. These windows of opportunity (or shocks) are critical events that, for a short period of time, break down the relevant barriers to change.

2.2.1 Incrementalism Applied to Technical Systems

Although defined in the policy context, the notion of planning large changes as a sequence of small, incremental improvements applies to technical systems as well. The basic idea is that stakeholders will be more tolerant to uncertainty in returns, if their upfront cost is either small or spread out over time. Thus, if the technical deployment can be layered such

that the cost and disruption associated with each change is small, resistance to the change will be less. This strategy is more or less feasible, depending on the nature of the technology in question. Scholars of changeability and evolvability have provided guidelines for how particular kinds of systems can be designed to be more changeable [Fricke and Shultz, 2005; Ross et al. 2008]. Similarly the real options framework provides strategies for designing systems that can be built-out later depending on how the use context evolves [Hassan et al. 2005]. The ability to leverage these ideas merits consideration during the design phase.

In the FAA case, it was found that a relatively minor redesign of the technical system could allow incremental adoption of the capability. This served to mitigate stakeholder risk due to free riding [Jenkins, 2009]. Currently, the military is considering several levels of capability change in the next generation ATN. However, each capability level is currently being considered as it's own end state. Technical staging of these capability levels may be an avenue for future consideration.

2.2.2 Leveraging Shocks to Catalyze Transition

Shocks have the effect of temporarily reducing stakeholders' resistance to change. However the longevity of the change window varies with the magnitude and intensity of the particular shock [Kingdon, 1984]. For example, in the air transportation context, newsworthy accidents, though tragic, often enable needed safety improvements [Mozdzanowska et al., 2008], but only when the fix is ready and can be deployed quickly. Unfortunately, history has shown that shocks in the form of imminent capacity constraints rarely have the same effect. In other words, travel delays, while annoying, don't tend to prompt airlines to invest in needed upgrades immediately.

As noted above, transition in path-dependent infrastructure systems often takes years to complete,[5] simply as a function of the distributed operational nature of the legacy system and the cost of the implementation. Thus, a single shock is rarely capable of sustaining buy-in long enough to complete a full upgrade of a system in its entirety. As a result, research on transition strategies has suggested leveraging shocks to enact smaller self contained changes [Campos et al 2007]. The idea being that adoption in a confined region of the system will prove-out potential benefits of adoption and in so doing encourage the rest of the system to adopt.

Several examples of this strategy can be seen in the National Airspace transition to NextGen. The FAA has achieved some success identifying self-contained early adopters. To date, they have executed regional programs in Alaska, the Gulf of Mexico and Louisville. In Alaska all weather visibility was the driving issue. With adoption of the new technology, General Aviation (GA) pilots were able to fly in bad weather with much lower risks. They saw clear, immediate safety benefits associated with better situational awareness [Campos et al., 2007] and could thus justify the upfront investment. This pilot program served to

---

[5] For example, the technology for Link-16, the current tactical datalink, was developed in the 1960s and it was initially deployed on some aircraft in the 1970s. Deployment of Link-16 has continued to the present day. It is worthy to note that there are still aircraft in the U.S. inventory that do not yet have Link-16.

prove the safety benefit to other GA pilots. In the Gulf, helicopter pilots were also enticed by the new ability to fly in bad weather. For them it was a direct financial benefit associated with predictable trips to oilrigs [Campos et al., 2007]. A slightly different logic applies to the pilot program in Louisville, where FedEx was given sole access to the airport overnight. Since as a single stakeholder they could guarantee full equipage, they were able to test out the promised new operating procedures. Thus, in a local setting, the FAA was able to quickly demonstrate the value of the upgraded fleet.

These examples illustrate that mini-shocks (inducing small-scale rapid change) can be manufactured when appropriate groups, regions, or platforms are targeted. The challenge is in identifying appropriate starting points, which can also lead to broader adoption by other groups, regions, or platforms. In the military tactical network context, analogous dimensions are starting with a single Service (e.g., Air Force vs. Navy), unified combatant commands (COCOM)[6] (e.g., specifically a regional command such as PACOM, AFRICOM, etc.), or aircraft type (e.g., Fighters vs. Intelligence, Surveillance and Reconnaissance (ISR) aircraft). In the sections that follow, we will explore how the appropriateness of a particular starting point is correlated to identifiable structural attributes of the particular system.

2.3 Alternative Transition Pathways: ATN Context

The above discussion provides a basis for defining the set of alternatives transition strategies that will be evaluated and compared in the remainder of the paper. Five transition strategies were chosen for analysis. They cover the range of mini-shocks feasible in this ATN context: one stakeholder first, one region first, and one platform first. Given the stylized nature of the model parameters, two bounding cases were also defined: the no-upgrades baseline, and the best-case scenario of everyone begins upgrading immediately. The seven total scenarios are summarized in Table I.

Table I: Summary of Seven Scenarios goes approximately here

3. Research Approach

In order to explore the relative value of the alternative technology transition strategies described above, we conducted a model-based case study of possible approaches to upgrading to the ATN. We chose the tactical network transition as the central case study because it is a critical example of lock-in associated with a path-dependent infrastructure system [Yin, 2009]. The core dynamics of the system are captured in a simplified model described below. Representative fleet characteristics were collected from open sources

---

[6] A unified combatant command (COCOM) is a "command with a broad continuing mission under a single commander and composed of significant assigned components of two or more Military Departments that is established and so designated by the President, through the Secretary of Defense with the advice and assistance of the Chairman of the Joint Chiefs of Staff" [Joint Staff, 2013]. There are six regional COCOMs and three functional COCOMs. Most forces are permanently assigned to continental United States (CONUS), but are apportioned out to the COCOMs as necessary for mission execution. Pacific Command (PACOM) and European Command (EUCOM) have forces assigned permanently to them. They, too, will draw on aircraft from the CONUS in the case of a major contingency.

[IISS, 2014]. Normal operating procedures, and associated variability, were abstracted from open sources, as well as informal interviews with current aircraft maintenance officers, current and former operational aircrew, and aircraft industry representatives.[7]

3.1 Model Overview

Our intention was to develop a basis for exploring relative differences among transition strategies and not to build a prescriptive/accurate model. To do this, we developed a model comprised of two main components: network capability (overall fleet characteristics) and network demand (representative scenarios). We also developed a range of possible network upgrade options. The interaction of the two components allowed us to generate usage statistics and analyze the impacts of different technology transition strategies. Once the overall fleet characteristics were determined for each potential approach using the first part of the model, the second part of the model applied those characteristics to a series of high-level missions to define what capability is delivered within each pathway. A conceptual overview of the model, and each of its component pieces, is shown in Figure 1. Each component of the model and associated data sources are described in the sections that follow.

Figure 1: Conceptual Overview of Simulation Model.

3.1.1 Network Capability: Evolution of Fleet Characteristics

The basic logic of the simulation follows the lifecycle of, and potential network upgrades to, each airframe. That logic is then aggregated across the fleet. The state of the fleet is tracked as a matrix of each individual airframe evolving over time. Each airframe is defined by six attributes: its military service association (i.e., Air Force, Navy), platform type[8] (i.e., fighter, bomber, ISR), its disposition (i.e., where it is in its cycle), the number of flight hours it has accrued, its upgrade level (defined in the next subsection), and the number of hours since its last maintenance. These attributes are outlined in Table II. Each attribute is updated at each time step. For the purpose of this simulation, a fleet of 2,500 was defined using only Air Force and Navy platforms. This provided sufficient diversity to examine relevant trends.

Table II: Fleet Characteristics goes approximately here

The model assumes that all aircraft are starting with a legacy network. Over the course of the simulation, airframes receive a single, backward compatible network upgrade. The timing of the upgrade is determined by the underlying constraints of when particular upgrades can be implemented (in terms of the disposition and maintenance cycles) and the

[7] Maintenance officers represented three broad aircraft categories; fighters, airlift and rotary wing. All were current, Active Duty Air Force officers with over 15 years of experience in their specialty. Aircrew represented experience in the Active Duty Air Force, the Air Force Reserve and the Air Guard as well as Special Operations Command. Experience ranged from 8-25 years, and over five different aircraft types. An industry representative with experience in operational flight program upgrades was also consulted.

[8] The U.S. DOD employs five additional aircraft classes (including command and control, transport, tankers, helicopters, and UAVs). The simplification of using only three assets types, i.e. fighter, bomber, and ISR, supports tractability without losing an ability to differentiate types of aircraft interactions. Note that for our purposes, we define ISR platforms as aircraft whose primary mission is to collect tactical ISR.

transition strategy in effect. We consider the underlying constraints to be true in all scenarios, where only one transition strategy (defined in Section 2.3) can be in effect for a given simulation. After all, the relative merit of the pathways is the outcome we are interested in.

Before describing how the disposition and maintenance modes are implemented, a brief note on how upgrades were integrated and assigned is required.

3.1.1.1 Technical Capability/Upgrade Levels: Individual Airframes

The network upgrade can take one of five levels, as defined in Table III. For this analysis we assumed a notional network upgrade option that provides order of magnitude capability increments. Correspondingly, the different upgrade levels require progressively more substantial changes to the existing hardware, and must be performed in different maintenance locations. The model assumes a single upgrade per airframe type because upgrades are both time-consuming and expensive. Nominal mission profile, age of the aircraft, and physical constraints of each airframe type determined the upgrade level. So for example, the RC-135, which will be recapitalized onto a new airframe in the coming years, receives only a limited upgrade. Similarly, the F/A-18B, which is being phased out of the fleet receives only a limited upgrade, while the F/A-18E, which is just coming into the fleet receives an Enhanced + upgrade. Additional details about model implementation can be found in [Rohrbach, 2013].

Table III: Upgrade Level Descriptions goes approximately here

3.1.1.2 Aircraft Disposition and Assignment

At each time step, airframes are assigned to one of four possible activities: Deployment, Training, Squadron Maintenance, and Depot Maintenance. The assignment is based on deterministic deployment cycles,[9] with variability around the particular timing of squadron-level maintenance. In the model, both mission and training flights accrue flight hours. Table IV summarizes the assumptions used for flight hour accrual for both training and deployment periods. Note that the deployed function is roughly three times the tempo of the home station tempo.

Table IV: Flight Hour Assumptions goes approximately here

A more detailed description of the mission assumptions during deployment is discussed in Section 3.1.2. The windows for maintenance arise stochastically as an aircraft continues to accrue flight hours. While the maintenance cycles occur after a deterministic number of operational hours (Table V), the time between maintenance cycles is stochastic because of the mechanisms through which flight hours accrue. The disposition module randomly selects aircraft for deployment from those available. As a result particular airframes accrue operating hours at randomly varying rates. After each deployment cycle, some eligible aircraft are sent back to depot for more substantial maintenance.

[9] Nominally, the Air Force follows a 20-month deployment cycle, spending 14 months in training, 2 months in deployment preparation, and 4 months deployed. For the Navy, the cycle is 24 months, including 14 months in training, 4 months in deployment preparation, and 6 months deployed.

Assuming an upgrade is scheduled, the model considers three methods for upgrading military aircraft: Time Compliance Technical Orders (TCTO), Modernization, and Block Upgrades.

1. TCTOs are upgrades that are made to a given platform type in the field by the squadrons as part of the regular maintenance process. Generally, TCTOs include swapping out one system or part for another. TCTOs are very common, but can take years to fully implement.
2. Modernization takes place at the depot as part of the regular maintenance cycle, and generally inserts a more modern system that is already in use in more advanced models of the same airframe. Usually, modernization results in greater equipment consistency across the fleet, as well as improved capability. It can include, for example, swapping out an engine or upgrading a radar system.
3. Block Upgrades are planned capability improvements for new models of a given airframe. The upgrades are built into new aircraft *on the production line*. Block or Increment Upgrades are generally much more extensive improvements, and usually include updates to the aircraft's Operational Flight Program (OFP). A new Block or Increment Upgrade occurs on average every 5–10 years [GAO, 2012].[10]

As defined in Table III, the first three levels of upgrade (Limited, Limited +, and Enhanced) can be performed during routine squadron maintenance as a TCTO. In the model, when particular aircraft are assigned to squadron maintenance, they are checked against the State of the Fleet matrix to assess whether an upgrade is required (i.e. there is a TCTO open for that specific airframe). If an upgrade is required, it is made at that time. Since squadron maintenance is typically quick, the aircraft is returned to duty in the next time step.

As part of the regular maintenance cycle, aircraft are sent back to the depot for in-depth checkouts, refurbishment, and more intrusive maintenance. In the model, when particular aircraft are assigned to depot maintenance, according to the timeline assumptions in Table V, their upgrade level and platform type are checked against the current state of the fleet portion of the model to assess whether an upgrade is required. Unlike in squadron maintenance, depot maintenance usually takes the aircraft out of service for some prescribed time. A limited number of any given aircraft type are in depot maintenance at all times.[11] As defined in Table III, Enhanced+ can be performed during depot maintenance.

Table IV: Timelines for Maintenance goes approximately here

Finally, the full implementation occurs only during Block Upgrades, when brand new aircraft are being produced. This category of upgrade was not implemented in the current model, since there is too much uncertainty associated with when future airframes will come online to make any outcomes realistic.

---

[10] Sometimes depot maintenance is used to implement Block Upgrades in aircraft with very long service lives or in cases where the production line for the airframe is no longer open (e.g. the Boeing 707 used by JSTARS). We do not account for these cases in our model, but in the real world these cases add a significant amount of time to the normal depot cycle.

[11] Per the inputs that we received from aircraft maintainers, roughly 10% of a given aircraft type could be expected to be in depot at any given point in time.

3.1.1.3 Network Capacity: Aggregated Capability

The current tactical network capacity is insufficient for the growing number of joint and coalition users and lacks flexibility and capacity / throughput [Smith, 2005]. One core purpose of the network upgrade is to alleviate these known capacity constraints.[12] On an individual basis, each level of upgrade is assumed to provide the potential for an order of magnitude increase in capacity to the network. Figure 2 represents each upgrade's capacity as a square to help visualize the differences. However, that full benefit is only realized when two aircraft of the same upgraded level are communicating with one another. In cases of mixed fleet operations (i.e., when aircraft of different upgrade levels operate together) all communications must occur at a level equal to the lowest capability in the link. For example, in the case of a group of five aircraft (10 possible pair-wise links total), where three are at a higher upgrade level than the remaining two, 70% (7 of 10 total links) of the mission must remain on the lower network, and 30% (3 of 10 total links) can move over to the upgraded network. Thus, the effective size of the upgraded network will expand and contract depending on the state of the aircraft using it; and "network effects" will necessarily create a lag between number of upgrades and benefits accrued.

Figure 2 goes approximately here: Conceptual representation of upgraded network usage

3.1.2 Network Demand: Representative Usage Scenarios

The second part of the model considers how fleet operations create demand for network capacity. The model defines three generic mission types (Surveillance, Close Air Support (CAS), and Strike)[13] in order to develop a representative understanding of the network usage over time. Assumptions about number and type of aircraft involved in each mission were vetted with current operators and are in Table VI. Our intention was not to create a realistic operational model, but rather to provide a basis for comparing relative system performance under the different scenarios.

Table VI goes approximately here: Mission Assumptions

The total network demand associated with a particular time period is a function of (1) the deterministic baseline demand profile which determines the network usage of each particular mission type and (2) random operational tempo patterns generated by the model which determine the number of missions flown in a particular period.

The baseline growth function was abstracted from several recent studies that project dramatic growth in the demand for communication network capacity of the next decade

---

[12] From the perspective of operational effectiveness, increased capacity generically can be considered a proxy for both better overall connectivity / interoperability (more network participants) and the use of more cooperative, and network intensive, applications.

[13] Aircraft are assigned to missions based on platform type and disposition (e.g., that the aircraft is available for deployment in the timestep and is an aircraft used for that mission). For Strike missions multiple types of fighter aircraft are used simultaneously to fulfill different aspects of the mission. This underscores the need to upgrade platform types across the fleet, as they will be fighting together in heterogeneous platform groupings. Strike missions also include ISR and bomber platforms. For the CAS mission, we assume two like fighter platforms are assigned to each individual mission, although all fighter platforms can be drawn on to perform this mission. ISR platforms are also part of the CAS mission.

[Haven, 2003; Furstenberg, 2012; Leland and Porche, 2004]. To account for significant variability in the projections, in lieu of running an extensive sensitivity analysis on an already highly variable model, we defined three future demand growth profiles that spanned realistic futures (e.g., low, moderate, and high). The mid-level expected growth is drawn directly from [Furstenberg, 2012]. Specifically, the data in that paper was fit with an exponential growth curve:

$$y = ae^{k(t-1)} \qquad (1)$$

Where a was found to be 0.95 and k was found to be 0.24. Since the purpose of the high and low baselines were intended to provide bounding scenarios, k was varied +/- 1. The three levels of baseline demand are depicted in Figure 3. The demand growth profiles are used to scale the network usage of each particular mission type and are updated on an annual basis. By using the three demand growth curves, different transition strategies can be compared in terms of their sensitivity to network demand.[14]

Figure 3 goes approximately here: Comparison of upgrade levels and demand profiles.

The stochastic part of the usage simulation determines the intensity of operations in a given time frame. In other words, it determines how many missions of each type are flown in a particular period. Historical U.S. military operations/engagement levels were used as the basis for intensity level assumptions [Leland and Porche, 2004]. [15]

Aircraft are assigned to missions based on platform type and disposition (e.g., that aircraft is available for deployment in the time step and is an aircraft used for that mission). Upgrade level is not considered in the assignment since it is a not considered in the real-world assignment decision.

For aircraft assigned to training, we use a fixed number of flying hours, based on inputs we received from active duty operators. These numbers vary by aircraft type, but are static across all demand profiles (see Table IV).

3.1.3 Model Output Measures

The main outcome of interest from this simulated fleet is the temporal relationship between the costs of implementing differing upgrade strategies and benefits to both

[14] That demand is increasing in not in doubt. The office the DoD Chief Information Officer has published a strategic plan to address these issues, but even the plan recognizes that "… to achieve interoperable infrastructure and synchronized operations, DoD must persuasively demonstrate that these strategies will improve computing and communication, and especially *provide the capacity to meet surge demand*."

[15] Low tempos require 75% of the current normalized network resources, while moderate and high tempos require 95% and 155%, respectively. Using history as a guide, the model assumes a major military contingency occurs approximately once every 5 years. During the course of a major contingency, critical, high tempo operations take place 6-10 times per year. Outside of a major contingency, the frequency of high tempo operations falls to approximately 2-5 per year. Each high tempo operation lasts for approximately 1-6 weeks and is randomly selected as the operational tempo is defined. The remainder of the time is broken up into moderate and low usage scenarios. For years where the U.S. is not involved in a major contingency, the model assumes that 60% of the remaining time is spent in a low operational tempo, and 40% is spent in the moderate tempo. For years where the U.S. is involved in a major contingency, 30% of the remaining time is spent in a low operational tempo, and 70% of the remaining time is spent at the moderate tempo.

particular stakeholders and stakeholders in the aggregate. Relative costs are accounted for directly as part of the simulation, assuming a one-to-one correlation between each upgraded aircraft and cost. Stakeholder benefits are related to the availability of required capacity to support planned operations. To that end, we tracked three types of outputs to summarize and compare the performance of different upgrade strategies: raw network overages, improvements attributable to upgrades, and proxy costs.

3.1.3.1 Network Usage

Real world network calculations are quite complex and are beyond the scope of this effort.[16] For our purposes, we focus exclusively on availability of network capacity as a proxy for benefit. The model measures the difference between doing nothing, which results in serious connectivity and capacity issues [Smith, 2005, DoD CIO, 2010-2012], and the various transition rollout options. Network usage at a particular time step is counted as a percentage of capacity used compared to capacity available.

In this context we conceptualized available capacity as a two-dimensional rectangle that is filled in (i.e., used) when missions are flown. The size of the available network expands and contracts with the technical capabilities of the airframes currently in use. The demand for network capacity is a function of the number and kind of missions being flown.

As defined in Section 3.1.2, the number of simultaneous missions are defined by the stochastic intensity component of demand, while the capacity required by a particular mission is defined relative to the original network (i.e., a typical surveillance mission occupies 25% of the legacy network) and grows exponentially according to the baseline demand profile.

As outlined in Section 3.1.1.3, the technical upgrades to an aircraft's communication system have the effect of expanding the available network size (i.e., the size of the two-dimensional rectangle), when communication is among upgraded aircraft. Therefore, pairs of aircraft communicate at the lowest level of their joint capability. The network capacity is calculated in proportion to the number of links at a particular capability level. The growth in capability thus inherently lags the number of upgrades. Per Figure 2, the capacity differences between upgrade levels are assumed to be one order of magnitude.

3.1.3.2 Raw Overages

Raw overages are directly related to network usage described above. They are a simple count of the number of times network capacity is exceeded by demand in a given six-month period. Since this is a stochastic model, different runs yield different specific values. The distributional nature of the results is captured, by plotting separate box-and-whisker plots for each time interval. While the distributions are large, trends are still apparent. Figure 4 illustrates how the data is represented, contrasting the moderate demand outcome for the two extreme (bounding) cases, everyone upgrades and no upgrades. A qualitative

---

[16] Creating a model of actual demand would include such factors as specific information exchange requirements; who "talks" to whom, packet size of the data being exchanged, throughput, duration, periodicity, latency requirements, distance between nodes (i.e., is routing required?), and acceptable jitter. We elected to focus on capacity.

description of the results for all cases is captured in Table VII, and discussed in detail in Section4.4. For full results, please see [Rohrbach, 2014].

Figure 4 goes approximately here: Raw overage measures for the Air Force assuming a) no upgrades (top) and b) everyone begins upgrading immediately (bottom).

3.1.3.3 Improvement attributable to upgrades

Improvements attributable to upgrades are calculated in comparison to the "no upgrade" case. In other words, we ask the question, how much better is performance with the upgrade than it would have been without? Figure 5a illustrates this notion for the "everyone upgrades" case. Initially, the upgrade does not fully compensate for the capacity constraints, but by year six, the capacity constraint is fully alleviated. Only mean values are shown.

Figure 5 goes approximately here: Everyone Upgrades Case a) Improvements due to upgrade (blue is Air Force, green is Navy); b) Costs of upgrade, measured as number of upgrades at each of Limited/L+ (blue) and Enhanced/E+ (red).

3.1.3.4 Proxy Costs

Actual cost data associated with military aircraft is tightly held, and was therefore not available for use in our analysis.[17] As a proxy, we tracked the number of upgrades in given time periods. In practice, the cost of upgrade integration will be quite substantial and vary significantly between upgrade levels (i.e., Limited vs. Enhanced). Rather than aggregating based on an assumed cost, different Limited/Limited+ and Enhanced/Enhanced+ upgrades were tracked separately to enable future aggregation, should better cost data become available. Figure 5b, shows the proxy cost plot for the "everyone upgrades" case. The upgrade numbers are shown on a per-period basis; all upgrades are completed by year eight. A government entity with access to actual cost data could use this outcome to calculate actual costs in dollar terms.

3.2 Model Validation

Because of the multiple levels of complexity within the model, it is a difficult model to validate. At each step of the development path, we not only used inputs and assumptions from experts, but we also took care to review model results and behavior with them to ensure that the dynamics of each portion of the model were consistent with the trends an experienced practitioner would expect. This process was conducted at each stage as the model was developed; initially experts validated the specific component they contributed to, but as the model was constructed, particular modules were validated as well. For example, the output of the fleet upgrade timeline was presented and discussed. While this particular upgrade hasn't been executed, the general diffusion timeline was checked. This resulted in several reworks along the way, and lively discussions about assumptions. We

---

[17] Cost data is *not* publically releasable at best, and is classified in many cases. As part of the work we had multiple discussions with the sponsor about the qualitative nature of projected costs. These discussions were the basis of our assumptions. Broadly, network upgrades are very expensive to implement, and are in some cases cost prohibitive. That is why we make the assumption that there is only one upgrade per airframe. While the model is simple, it does allow us to explore the impact of different transition strategies, and our results show that there are different outcomes.

feel confident that the model is a decent, first order approximation of the system. Although there is no doubt room for improvement in the model assumptions, the representative output plots resulting from the model illustrate important trends and demonstrate the value of exploring transition paths in this way.

3.3 Model Runs

Each of the seven transition strategies (Section 2.3) was explored over a 15-year period across the three possible future demand scenarios described in Section 3.1.2, for a total of 21 cases, simulated in the model described above. One thousand runs were conducted for each of the 21 cases to produce stable outcome distributions. The results are described in Section 4.

4. Case Study Results

This section summarizes the case-specific results for each of the upgrade strategy alternatives. The first and last strategies were run first to establish boundary conditions. They were run against all three of the demand profiles. If there are no upgrades, and there is high exponential growth in demand, the legacy network will be exceeded 100% of the time within three years. Under both low and moderate exponential growth, the legacy network will be exceeded 50% or more of the time within three years. If everyone upgrades simultaneously, there will be substantial surplus in capacity under realistic conditions (i.e., in all the cases simulated) 5 ½ years after upgrading begins. The variants of each of the platform, service, and regional upgrades were explored next.

4.1 Platform Increments

In the context of our simplified fleet, we explored the potential benefits of prioritizing upgrades to one particular type of platform within the fleet. We selected ISR platforms because the number of ISR platforms is small, and therefore the upgrade would be relatively inexpensive. These upgrades were then followed by upgrades to other platform types. Upgrading ISR first provides consistent benefit (i.e. lower percentage of shortfalls as compared to the no upgrade case), but the benefit is relatively small. We compare that to the opposite strategy of prioritizing fighters first. Upgrading fighters first provided high benefit only in the low demand case. Given that most missions use some combination of fighters, bombers and ISR platforms, no concentration of local benefits is achieved when only one platform type is upgraded; the main selling point is that by limiting the number of platforms receiving an upgrade, the initial cost is relatively lower than the cost of upgrading the entire fleet.

In more general terms, platform-first does not create a mini-shock because platforms do not operate independently by type, and can therefore not achieve critical mass, more quickly, on their own.

4.2 Stakeholder Increment

The DoD has often used stakeholder-specific increments historically[18] because such strategies drastically simplify challenges associated with distributed decision authority. A Navy-first strategy was explored in the model because the Navy does sometimes execute missions independently; especially in low demand scenarios. Pursuing a Navy-first strategy therefore gives us an opportunity to explore the full benefits of an upgrade within a small group of actors. While this approach merits detailed consideration if stakeholders are misaligned, it is otherwise of limited value. While the Navy does operate independently in some cases, most DoD operations, especially in the case of a major contingency (e.g. the high demand scenario), integrate multi-service capabilities and platforms. Therefore, there are limited opportunities for a single service to achieve critical mass independently. As a result, the non-adopting service (in this case the Air Force) will receive free benefits, as the Navy increasingly moves parts of their operations to the upgraded network, but will not see the full deployment benefits that might encourage them to adopt themselves.

An additional surprising insight that was revealed through the analysis, is the inherent catch-up effect associated with delaying investment. This phenomenon is observable because we imposed the delay but is not related to the particular scenario. It turns out that when exponential growth is assumed (even in the weak growth “low” scenario), waiting even a few years to begin the technology transition will lead to a longer delay in benefits. This phenomenon is illustrated in Figure 6.[19] In this scenario, the Navy begins investing in year 0 and starts seeing benefits by year 2 – a 2-year lag in benefit. Since we are modeling a situation where the Air Force waits to see Navy benefit before beginning to deploy the new capability, their initial investment isn’t made until year 5. As shown in the figure, the Air Force does not see much in the way of benefits until year 10 – a 5-year lag in benefit. This additional 3-year delay occurs because, by year 5, the Air Forces capability has already degraded. As a result, although the investment in new capability improves the situation as it does with the Navy, there is a much longer way to go before benefits are observable.

This is important more generally because lags in benefit accrual have been shown to be the difference between maintaining stakeholder buy-in or not. In this case, the notion of a Navy-first strategy was intended to mitigate stakeholder resistance. However, this observed catch-up effect may limit the effectiveness of such a strategy.

Figure 6 goes approximately here: Illustration of catch-up effect when Navy upgrades first.

---

[18] For example, Cooperative Engagement Capability (CEC), which was deployed within the Navy by Carrier Task Force, built momentum and user pull because the full capability was demonstrated.

[19] The differences in the rate of benefit accrual between Navy and UASF are artifacts of the way the services operate their fleets in the simulation. In some of the mission scenarios considered, the immediate upgrade of even two Navy airframes can yield a noticeable step-change benefit to the network (in the aggregate it shows up as a 20% starting benefit and 50% by year 3). The Air Force on the other hand is involved in more mixed-fleet operations and missions tend to involve more airframes. As a result, the benefit shows a more steady growth, as network effects take hold in the aggregate. The basic point of Figure 6 is the difference in lag.

4.3 Regional Increments

There are six regional unified combatant commands (COCOMs) and three functional COCOMs. Most forces are permanently assigned to the continental United States (CONUS), but are apportioned out to the regional COCOMs as necessary for mission execution. Pacific Command (PACOM) and European Command (EUCOM) have forces assigned permanently to them. They, too, will draw on aircraft from the CONUS in the case of a major contingency.

In the context of our simplified model, only a two-region system is considered: a primary region, representative of CONUS, and a COCOM with its own assigned forces (e.g., PACOM or EUCOM). This allows us to consider the interaction among regions during high tempo operations, but does not provide a basis for looking for differences among regions. In terms of upgrade strategies, both a COCOM-first and CONUS-first upgrade is explored. The benefits to each are evaluated from the perspective of both regions (i.e., COCOM-first as viewed from the COCOM, COCOM-first as viewed from CONUS, CONUS-first as viewed from CONUS, and CONUS-first as viewed from the COCOM).

4.3.1 COCOM Upgrade

Viewed from the upgraded COCOM: Since the number of aircraft permanently assigned to a COCOM is relatively small, the entire upgrade can be performed fairly quickly. We expected this to mirror the "everyone upgrades" case. What we found was that in the low demand scenario, this approach provides a significant and immediate improvement, matching the "everyone upgrades" case for the specific region. However, when demand increases beyond the self-sufficiency of the COCOM, aircraft will begin to be drawn from CONUS (which has not yet begun to upgrade). As a result, in high tempo operations, the COCOM continues to experience network shortfalls.

It should be noted that much of the impact of receiving non-upgraded assets into the theater depends on how those forces are used. If they are fully integrated into operations, as assumed here, there will be a negative impact on the network. If, however, CONUS based forces are used as backfill for COCOM forces, which is the current method of operation, the negative impact may not be as pronounced as seen in the results.

Viewed from the non-upgraded CONUS: Identical to the no upgrade scenario within CONUS, since CONUS never pulls aircraft from other regions. When deployed to the upgraded COCOM, CONUS forces will benefit from additional network resources made available as the COCOM forces increasingly move parts of their operations to the upgraded network.

4.3.2 CONUS Upgrade

Viewed from CONUS: A CONUS-specific upgrade would look much like an everyone-all-at-once upgrade, as viewed from CONUS, since CONUS is self-sufficient. However, recognizing that the United States would have to be under attack, for engagements rarely occur in the CONUS region, and that the fleet is large, the CONUS-centric value is not in-and-of-itself appealing. Additionally, there is a possible negative outcome. Most training occurs in the CONUS. This means that the capability available during training will be significantly better than the actual combat capability. So it is likely that the capability will be decremented for training to reflect the actual combat capability, negating the improvements of the upgrades.

Viewed from a non-upgraded COCOM: A somewhat surprising result of the simulation is the second-order benefit that a CONUS upgrade provides to a non-upgraded COCOM. Since COCOMs pull CONUS aircraft when operations are intense, a CONUS upgrade would in fact delay the onset of major overages in a COCOM for several additional years, assuming CONUS-based forces are fully integrated into operations. However, those improvements will be limited if CONUS-based forces are used only as backfill capability (current method). These dynamics are shown in Figure 7 – although we would nominally expect no benefit in the COCOM (since it has not undergone an upgrade), an initial benefit is observed. This is a second order benefit due to the CONUS upgrade (i.e., COCOM is pulling upgraded benefits from CONUS). The value is not substantial, but it is an option that may merit follow-on study. We expect that if implemented, a COCOM follow-on upgrade, not shown in the current data, would be planned (leading to improved performance in the out years, compared to what is shown).

Figure 7 goes approximately here: Second-order improvements experienced in a COCOM, due to a CONUS upgrade.

4.4 Overview of results

Table VII provides an overview of the key takeaways across the cases.

Table VII goes approximately here: Overview of Key Takeaways Across Cases

5. Discussion

Having discussed the model-specific results above, this section discusses the implications for future technology transitions. Section 5.1 focuses on specific takeaways relating to the central case. Section 5.2 explores more general transition principles.

5.1 Case-Specific Insights

Given the results presented in Section 4, what guidance can be provided to a systems architect tasked with designing a transition from the legacy network to the upgraded IP-based system described above?

First, it must be recognized that even if a full-fleet upgrade is initiated immediately, the process will take a minimum of six years. Of the assumptions used in the model, those associated with the maintenance cycle are the most grounded in current Air Force and Navy practice. We assumed varying levels of upgrade that could be accomplished from work done by squadron level maintenance, all the way up to depot maintenance. Low-level upgrades at the squadron level can be completed in an approximately 2–2.5 year timeframe. Upgrades at the depot are completed in 4–6 years, depending on the programmed depot maintenance schedule. As a result of this extended period of mixed fleet operations, backwards compatibility of the technical system is critical given that the aircraft must be able to work together effectively during the six year upgrade period. Backwards compatibility had not been an explicit assumption, but given the amount of time a projected rollout would take it is certainly a requirement that would be worth associated additional costs.

Second, the transition strategy can have a significant impact on how costs and benefits are spread across the system over time. Given the integrated nature of nominal Air Force and

Navy operations, we found limited value in platform- and service-specific incremental strategies. However, a more detailed exploration of region-specific increments merits further work. A moderate sized, relatively isolated, moderate operations tempo (OPSTEMPO) region might provide an ideal proving ground for the value of a network upgrade, creating a mini-shock capable of catalyzing a wider systems transition. The current model is not calibrated to realistic inter-regional dynamics or geographical operations. However, with appropriate classifications, this would be a straightforward extension of the current work.

Third, assuming that demand continues to grow (even weakly) exponentially, the sooner a particular upgrade implementation can start, (a) the delay will be shorter between upfront costs and future benefits (recall Figure 6), and (b) the longer that a particular upgrade will satisfy demand (recall Figure 3). If current demand projections are remotely accurate, no upgrade level will be adequate indefinitely. There is, therefore, an inherent tradeoff that must be made among (1) performance increments between generations of the infused technology, (2) time intervals between upgrades, and (3) continuity of service.

To illustrate this point, Figure 3 plots the three levels of projected demand, as well as each of the assumed network upgrade capability levels over the next 15 years.

- Limited/Limited+ upgrades can buy the military a decade of capacity relief, even under worst-case demand expectations. These upgrades can be easily implemented at the squadron level and could be infused in the very short term.
- In the near-term, Enhanced/Enhanced+ upgrades would provide capacity far in excess of what is needed. These upgrades would be effective for several more years beyond when Limited/Limited+ upgrades are overcome (much longer in the low demand scenario), but another upgrade would still be required by the two-decade mark.
- The full implementation of the network upgrade will not be needed for two decades, even in the worst demand case.

Future work should thus explore the potential for staged technical deployment, or at least planned technical evolution of the upgraded system. Recall the assumption that any particular aircraft will only get one upgrade (and that upgrade will be of a particular level), depending on that aircraft's expected longevity and technical integration characteristics. However, our results indicate that there may be substantial value to defining a transition strategy that leverages a staged technical deployment. If the various upgrade levels can be designed such that initial increments can be deployed in the field, and augmented when aircraft return to depot for planned upgrades, service can be maintained over time with much lower cost. This logic of explicitly designing for evolvability – be it stages of the current system, or facilitating future change out – is an extension of current thinking in the change literature [c.f. Fricke and Shultz, 2005] can have substantial cost implications for future upgrades.[20]

[20] Of course, realistic cost data is required to do this analysis.

5.2 Guidelines for Technology Transitions in General

While this study has focused on the implementation of one type of technology in one system context, we believe that our analysis yields insights into how technology transitions can be architected more effectively in path dependent, multi-stakeholder, infrastructure systems more generally. As with the literature review in Section 2, this discussion separates technical and institutional considerations. The concepts are then re-combined to support a framework for designing transition strategies.

5.2.1 Considerations in Architecting for System Evolution

Systems engineers are typically taught to minimize the need for downstream changes through careful upfront planning [Clark and Fujimoto, 1991]. The rationale is that changes are increasingly costly with every next design phase. However, in today's constantly changing environment, this viewpoint is no longer realistic; enforcing "closed" systems designs unnecessarily restricts technical evolution, and limits value delivery under changing contexts [Ross et al., 2008; Fricke and Schulz, 2005]. Clearly, not all systems should be infinitely evolvable. Steiner [1998] and Fricke and Shultz [2005] identified several defining characteristics for when changeability should be planned for. Grounded in their automotive context, they focused on production oriented new product development (NPD), where many products might be built from a basic set of core attributes, or have a stable core functionality coupled with variable secondary functions.

In our context, the related notions of changeability, evolvability and obsolescence management apply internally to a particular product. While the core building block notion is fundamentally the same, secondary functionality must be able to be added to the base system over time. For example, in the tactical network system described above, there are two ways in which evolvable design could be considered. First, with respect to backwards compatibility, given the constrained nature of the relevant hardware and depending on the nature of the new technology, enforcing a requirement for backwards compatibility on the new hardware may come at a substantial cost. It may instead be more appropriate to consider, for example, the limited upgrade to all legacy hardware, as a forward compatibility patch on the old network. Backwards compatibility for the new hardware would become a non-issue. Viewed this way, the whole fleet could be upgraded to baseline (limited) compatibility with a software patch, independent of the planned hardware swap-outs. Second, with respect to the upgrade rollout, a technically incremental rollout could mitigate many of the challenges discussed above. This would require the technical system to be designed as "building block" layers, infusible at different maintenance layers.

Viewed more broadly, any system that has a planned lifetime in excess of a decade and has a prescribed maintenance cycle governing multiple system components, should consider the interaction of its various natural time cycles with the technology being transitioned [Herald et al. 2009]. Given the inherent and variable diffusion lags, and continuous growth in demand, for systems like NextGen and ATN, incremental technical deployment can be a powerful rollout strategy. The key is to define intermediate states that will (a) deliver value at every point and (b) meet projected demand for the time between likely increments. The specifics of the technical solution should follow the kinds of modularity [Utterback and

Abernathy, 1975; Baldwin and Clark, 2000] and principles of changeability [Fricke and Schulz, 2005] and real options [Ricci et al, 2014] defined in previous work.

5.2.2 Considerations in Planning for System Evolution

Policy studies generally resign themselves to transitioning through either sub-optimal incrementalism or not-to-be-wished for crisis-induced shocks. In this work we explored the potential for planned mini-shocks as a means to garner stakeholder buy-in by limiting lags between investment and benefit accrual, and to demonstrate promised performance gains on a small scale. This notion has been explored theoretically in prior studies [Marais and Weigel, 2006] and implemented by the FAA, as part of NextGen. Our contribution is to systematically compare multiple dimensions of "mini-" in a simulated environment. While our model was not sophisticated enough to address social aspects of diffusion after the mini-shock, it does highlight several characteristics that make for quality mini-shock candidates.

Past studies have highlighted pressing need or obvious benefits as a core selection criterion [Marais and Weigel, 2006]. While this external attribute can facilitate buy-in for the mini-shock, it often limits the usefulness of the test as a demonstration – external parties may view the success as related to the "special-case" nature of the test (and therefore not representative of their own circumstances, e.g., Alaskan demonstration). To overcome this limitation, we considered representative sub-problems as a basis for mini-shocks. In particular, we focused on platform-, stakeholder-, and region-specific increments. Although each increment represented a logical sub-problem, we found that the structure of the overall system created substantial differences in appropriateness (from the perspective of mini-shocks). In particular, regional increments hold the most promise. This is because regions are loosely coupled from one another, while stakeholder groups, and platforms interact strongly within regions. In the specific example of ATN, there is the added advantage that many stakeholders would be exposed to regional benefits due to the rotational nature of personnel deployments. Given that in the ATN, lines of design authority can run in any of the three dimensions, focusing on decoupled sub-problems, with defined lines of authority is good advice for any transition.

5.2.3 Guidelines for Successful Transitions

This work has outlined two key dimensions that must be considered when designing a successful technology transition that involves the infusion of a technology into a networked system. First, the evolvability of the subsystem (e.g., the radio in the ATN example) must be considered first: specifically, whether the deployment of the system can be staged, such that increments of value are delivered through progressively invasive systems changes. Second, the structure of the network must be considered: specifically, whether there are decoupled sub-networks that can be exploited as mini-shocks. Depending on the potential to decompose deployment in either dimension, different transition paths are more or less feasible and appropriate. As outlined in Figure 8, if a mini-shock is possible, it should be exploited, with an emphasis on local demonstration of full, normal benefits. If options are available, single-stakeholder increments generally receive less resistance. If a technical increment is possible, it should be defined such that value is delivered at each increment, with repeat costs minimized. If possible, increments should proceed from less to more

invasive. Finally, if both network and system increments are possible, a full deployment mini-shock should be exploited first, but the technical increment option should be assessed. Depending on expected levels of resistance to the full deployment, the added cost of technical increments may not be worthwhile.

Figure 8 goes approximately here: Dimensions of problem decomposition in a technology-network system.

5.2.4 Future Work

In the context of the ATN, our research points to several areas that deserve additional consideration. We believe it would be fruitful to explore the following:

Actual cost data: Re-running the current analysis using actual cost data would confirm the specific results and recommendations to the ATN.

Staged technical deployment potential: We explicitly limited each aircraft to a single upgrade primarily because we did not have access to actual cost data. Our results show that the earlier the upgrade takes place, the longer the impact of the benefit. Relaxing the upgrade constraint to explore a staged technical deployment is a natural extension of the current work. Specifically, if two upgrades per aircraft are an option and the upgrade levels can be designed such that initial increments can be deployed in the field, and then augmented by a more extensive increment when aircraft return to depot for planned upgrades, it may be that benefits can be maintained over time with much lower cost.

Inter-regional dynamics: Our results show that region-specific upgrades hold the most promise. A more detailed model based on actual COCOM deployment statistics has the potential to tease out some of the dynamics that our initial results only hint at. A moderate sized, relatively isolated, moderate operations tempo (OPSTEMPO) region might provide an ideal proving ground for the value of a network upgrade, creating a mini-shock capable of catalyzing a wider systems transition. A model using actual historical deployments and current operational plans should enable a better understanding of the CONUS-COCOM dynamic, and lead to more specific and actionable staging recommendations.

Finally, more broadly, path-dependent systems remain an under-studied but increasingly important class of system change. Additional case studies are an important contribution to systems engineering research.

6. Conclusion

As engineering systems become more and more complex, technology transition increasingly involves deploying an upgraded subsystem across a legacy network. This mode of upgrade presents new challenges for systems architects concerned with maintaining value over multiple infused technical changes. This paper takes a systems perspective to explore the way that designing the transition path can influence system outcomes. It uses a model-based case study of planned upgrades to the ATN as a basis for identifying key attributes of component-network change.

With respect to the ATN case study, first, we identified a natural diffusion cycle of the system that was determined by the interaction of maintenance cycles and operating

procedures. Even if equipage started today, full deployment will not be achieved for six years; as a result, a mixed fleet operational strategy is critical. Second, some mini-shocks will be more cost effective than others. The analysis compared three strategies for staged-deployment (i.e., platform-, service- and region-specific increments) and found that the extent of decoupling the sub-network is an important determinant of the effectiveness of the strategy. In the ATN, region-specific upgrades hold the most promise and should be explored in more detail in future work. Third, we identified an inherent tradeoff between upgrade cycle and sustained capability levels. In other words, assuming even weakly exponential growth in demand, there is a relationship between timing of infusion and longevity of benefit. As a result, a less capable system, deployed expediently can do more good than a more sophisticated upgrade that can only be integrated in the next block upgrade.

More generally, this work identified two key dimensions of problem decomposition that form the basis for path design in technology transition. The first dimension is closely related to traditional systems engineering concern with changeability and looks at staging the technical deployment. The second speaks to attributes of the network structure, and seeks to identify decoupleable sub-networks that can be used to prove benefits quickly. When network effects are important, achieving local critical mass can encourage later adoption by other stakeholders. For example, the FAA has used this strategy of partial deployment, though their focus was on critical needs rather than normal isolation. The second dimension is more closely related to traditional systems engineering and looks at staging the technical deployment. Our core contribution is in integrating the two dimensions in the context of technology-network systems. Opportunities to stage should be explored in both dimensions in this class of system. Future work can develop specific guidelines for screening opportunities.

This class of technology-network systems is becoming increasingly important. Situational Awareness in the National Airspace, and mission-critical mobile communications that support military operations, are just two examples. While these networks have not traditionally been considered as systems to be architected, we believe that there is an important role for the tools developed to evolve physical products to play in this new setting. This work has taken a first step towards that goal.

References

ASD (NII)/DoD CIO. Department of Defense Information Enterprise Strategic Plan, 2010-2012. Accessed multiple times in 2014.

Adams, W. and Adams, W. J. 1972. “The Military-Industrial Complex: A Market Structure Analysis,” American Economic Review 62 (1/2): 279–287.

Anderson, P. and Tushman, M. L. 1990. “Technological Discontinuities and Dominant Designs: a Cyclical Model of Technological Change,” Administrative Science Quarterly (Adm. Sci. Q.) 35 (4): 604–633.

Baldwin, C. and Clark, K. 2000. Design Rules, Volume 1: The Power of Modularity. Cambridge US-MA: MIT Press.

Campos, N., Holcombe, R., Leybovich, M., Szajnfarber, Z., and Thorsteinsson, H. 2007. “Identifying the Dynamics of Technology Transition: ADS-B Adoption in the National Airspace System.” Final Report, Massachusetts Institute of Technology (Cambridge, US-MA).

Christensen, C. 2003. The Innovator’s Dilemma. New York US-NY: Harper Collins.

Clark, K. and Fujimoto, T. 1991. Product Development Performance, Strategy, Organization and Management in the World Auto Industries. Cambridge, US-MA: Harvard University Press.

Fitzgerald, M. E. and Ross, A. M. 2012. “Mitigating contextual uncertainties with valuable changeability analysis in the multi-epoch domain,” Presented at the Systems Conference (SysCon), Vancouver, BC, 19–22 March.

Fricke, E. and Schulz, A. P. 2005. “Design for Changeability (DfC): Principles to Enable Changes in Systems Throughout Their Entire Lifecycle.” Systems Engineering 8 (4) 342–359.

Furstenberg, D. 2012. “Intel: Meeting the Growing Bandwith Demands of a Modern Military,” MilsatMagazine, March.

Government Accountability Office. 2012. *Tactical Aircraft: Comparison of F-22A and Legacy Fighter Modernization Programs*, GAO-12-524, Washington, D.C., April 26, 2012.

Grimes, John. ASD/NII. Department of Defense INSTRUCTION NUMBER 4630.09, November 3, 2008.

Hassan, R., de Neufville, R., and McKinnon, D. 2005. “Value-at-risk analysis for real options in complex engineered systems,” Systems, Man and Cybernetics 4: 3697–3704.

Hastings, D.E. 2007. “Measuring the Value of Flexibility in Space Systems: A Six-Element Framework,” Systems Engineering 10 (1): 26–44.

Haven, J. 2003. “Transformational Communication,” presented at the Military Satellite Communications (MILSATCOM) Industry Days, El Segundo US-CA, 5–7 March.

Henderson, R. and Clark, K. 1990. “Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms,” Administrative Science Quarterly (Adm. Sci. Q.) 35 (1): 9–30.

Herald, T., Verma, D., Lubert, C and Cloutier, R. 2009 “An obsolescence management framework for system baseline evolution—Perspectives through the system life cycle.” Systems Engineering 12(1).

Institute for International Strategic Studies. 2014. The Military Balance. London, UK. International Institute for Strategic Studies.

Jenkins, M. R. 2009. “A systems approach to identifying aircraft equipage requirements, benefits and risks of ADS-B applications,” Masters Thesis, Massachusetts Institute of Technology (Cambridge, US-MA).

Joint Staff. Joint Publication 1: Doctrine for the Armed Forces of the United States, 25 March 2013, p. GL-11.

Kingdon, J. W. 1984. Agendas, Alternatives and Public Policies. Washington US-DC: Addison-Wesley Educational Publishing Inc.

Leland, J. and Porche, I. 2004. Future Army Bandwidth Needs and Capabilities. Santa Monica, US-CA: RAND Corporation.

Lindblom, C. 1992. "The science of muddling through," Public Administration Review 19 (2): 79–88.

Marais, K. and Weigel, A. 2006. "A Framework to Encourage Successful Technology Transition in Civil Aviation." Presented at the 25th Digital Avionics Systems Conference, Portland, US-OR, 15–19 October.

Mozdzanowska, A. L. 2008. "System Transition: Dynamics of Change in the US Air Transportation System." PhD diss., Massachusetts Institute of Technology (Cambridge, US-MA).

Mozdzanowska, A. L., Weibel, R. E., and Hansman, R. J. 2008 "Feedback Model of Air Transportation System Change: Implementation Challenges for Aviation Information Systems." Proceedings of the IEEE96 (12): 1976–1991.

Ricci, N., Rhodes, D. H., & Ross, A. M. 2014. "Evolvability-Related Options in Military Systems of Systems." Procedia Computer Science, 28, 314-321.

Richards, M. G., Ross, A. M., Shah, N. B., and Hastings, D. E. 2009. "Metrics for Evaluating Survivability in Dynamic Multi-Attribute Tradespace Exploration," Journal of Spacecraft and Rockets 46 (5) 1049–1064.

Rohrbach, A. K. 2013. "Implementing Technical Transitions through Schedule-Based Policy: Insights from the Military's Airborne Tactical Network," Masters Thesis, Massachusetts Institute of Technology (Cambridge, US-MA).

Rosen, S. P. 1994. Winning the Next War. Ithaca, US-NY: Cornell University Press.

Ross, A. M., Rhodes, D. H., and Hastings, D. E. 2008. "Defining changeability: Reconciling flexibility, adaptability, scalability, modifiability, and robustness for maintaining system lifecycle value." Systems Engineering 11 (3) 246–262.

Sapolsky, H. M. 1972. The Polaris System Development: Bureaucratic and Programmatic Success in Government. Cambridge, US-MA: Harvard University Press.

Scherer, F. M. 1964. The Weapons Acquisition Process: Economic Incentives. Cambridge US-MA: Harvard University Press.

Schumpeter, J. A. 1934. The Theory of Economic Development. Cambridge, US-MA: Harvard University Press.

Smith, Mark. 2005. "Link 16 Dynamic Network Management," presented at IDLS October 2005

Steiner, R. 1998. "Architecture and Evolvability – Definitions and Perspective," presented at the Proceedings of the 8th Annual Symposium of INCOSE, Vancouver, BC, 26–30 July.

Stigler, G. J. 1971."The Theory of Economic Regulation," The Bell Journal of Ecomonics and Management Science 2 (1): 3–21.

Szajnfarber, Z., Richards, M. G., and Weigel, A. L. 2011. "Challenges to Innovation in Government Space Agencies," The Defense Acquisition Research Journal (ARJ) 18 (3) Issue 59: 257–276.

Utterback, J. M. and Abernathy, W. J. 1975. "A dynamic model of process and product innovation," Omega 3 (6): 639–656.

Utterback, J. M. 1994. Mastering the Dynamics of Innovation. Cambridge, US-MA: Harvard University Press.

Yin, R. 2009. Case Study Research: Design and Methods – Applied Social Research Methods Series, Volume 5. Thousand Oaks, US-CA: Sage Publications, Inc.

Tables

Table I: Summary of Seven Scenarios

Strategy	Variants	Additional Comments
1 – No Upgrade		Low end bounding case.
Platform-Specific Upgrade	2 – Fighters and Bombers Upgrade First	A global upgrade, across only fighter and bomber platforms. Once these platforms are fully upgraded, ISR begin upgrading.
	3 – ISR Upgrades First	A global upgrade, across only ISR platforms. Once the upgrade is complete, Fighters and Bombers begin upgrading.
Stakeholder Upgrade	4 – Navy Upgrades First	Navy aircraft upgrade first, followed by Air Force aircraft.
Regional Upgrade (from COCOM perspective)	5 – CONUS Upgrades First	The region is not upgraded, so when assets are drawn from outside – i.e., from the CONUS – they will be upgraded, and therefore they increase the overall theater network capability.
	6 – COCOM Upgrades First	Aircraft in the region are upgraded first, so when assets are drawn from outside – which always happens for the high-demand scenarios – they will not be upgraded, and therefore they decrement the network capability.
7 – Everyone Upgrades Simultaneously		High end bounding case. A global upgrade, across all platform types. From the regional perspective this results in mixed upgrades across the OPSTEMPO mixes.

Table II: Fleet Characteristics

Attribute	Category Options	Static or Dynamic?
Military service association	Air Force Navy	Static
Platform type	Fighter Bomber ISR	Static
Disposition	Deployment Training Squadron maintenance Depot maintenance	Dynamic
Number of flight hours accrued	Formula based different assumptions for training versus operational flight hours	Dynamic
Upgrade level	Limited mode Limited+ Enhanced mode Enhanced+ Full implementation	Dynamic (note that a single, specific upgrade is assigned to each platform type)
Number of hours since last maintenance	Simple time calculation	Dynamic

Table III: Upgrade Level Descriptions

Mode	System change	Service required
Limited mode	Software upgrade only	Squadron maintenance
Limited+	Software and small hardware upgrade	Squadron maintenance
Enhanced mode	Moderate hardware and software upgrade	Squadron maintenance
Enhanced+	Full hardware and software upgrade	Depot maintenance
Full implementation	“Baked in” changes, includes new antenna	Block Upgrade

Table IV: Flight Hour Assumptions

	Home Station	Deployed
Bomber	2 aircraft per day x 8 hours per sortie	6 aircraft per day x 14 hours per sortie
Fighter[21]	8 aircraft per day x 1.3 hours per sortie	24 aircraft per day x 4 hours per sortie
ISR	2 aircraft per day x 4 hours per sortie	6 aircraft per day x 12 hours per sortie

[21] Fighter squadrons are generally comprised of 24 aircraft. On any given normal day, the execute a "12 turn 8." What this means is that of the 24 aircraft in the squadron, 12 will be flown in the first "go" of the day, and 8 of that original 12 will be flown in the 2nd "go." This means that 12 aircraft are in some combination of maintenance and configuration changes (e.g. mounting fuel tanks). But, they are generally not in maintenance for very long.

Table V: Timelines for Maintenance

Type	Squadron Maintenance	Depot Maintenance	% of fleet in Depot
Bombers	Every 90–120 hours, < 1 week (one timestep)	Every 5 years, for 7 months	10%
Fighters	Every 90–120 hours, <1 week (one timestep)	Every 5 years, for 7 months	10%
ISR	Every 270–360 hours, <1 week (one timestep)	Every 8 years, for 9 months	10%

Table VI: Mission Assumptions

Mission	Aircraft Assumed
Surveillance	2–5 ISR aircraft
Close Air Support	2 fighter aircraft
Strike	2 bomber aircraft 18–58 fighter aircraft (20–60 total)

Table VII: Overview of Key Takeaways Across Cases

Strategy	Low Exponential Growth	Moderate Exponential Growth	High Exponential Growth
No Upgrade	After year 2, shortfalls > 50% After year 12, exceeded 100%	After year 2, shortfalls > 50% After year 4 ½, exceeded 100%	After year 2, shortfalls > 50% After year 3 ½, exceeded 100%
Platform-Specific Upgrade (ISR First)	Through year 5, shortfalls up to 40% Years 6-9, shortfalls between 40-60% Year 9-10, shortfalls between 70-90% After year 10, shortfalls start to decrease	Through year 4, shortfalls up to 40% Years 6-9, shortfalls between 40-60% Year 9-10, shortfalls between 60-80% After year 10, shortfalls start to decrease	Through year 4, shortfalls up to 40% Years 6-9, shortfalls between 40-60% Year 9-10, shortfalls between 60-80% After year 10, shortfalls start to decrease
Platform-Specific Upgrade (Fighters First)	Through year 4 ½, shortfalls up to 50% After year 4 ½, no shortfalls	Through year 4, shortfalls up to 50% Through year 10, shortfalls 60-100% After year 10, shortfalls start to decrease	Through year 2 ½, shortfalls up to 50% Through year 10, shortfalls 70-100% After year 10, shortfalls start to decrease
Stakeholder Upgrade (Navy First)	For Navy through year 3, minor shortfalls, after year 3 has no shortfalls For Air Force through year 11 ,50% shortfalls	For Navy through year 4, minor shortfalls, after year 3 has no shortfalls For Air Force through year 5, 50% shortfalls Years 5-11, 80-100% shortfalls	For Navy through year 3, minor shortfalls, after year 3 has no shortfalls For Air Force through year 3, 50% shortfalls Years 3-11, 80-100% shortfalls
Regional Upgrade (COCOM Upgrades)	For CONUS, mirrors No Upgrade Case For COCOM, through year 5 shortfalls 10-20% Years 5-9 no shortfalls Years 9-13, shortfalls 15% After year 13, shortfalls over 50%	For CONUS, mirrors No Upgrade Case For COCOM, through year 7 shortfalls 10-30% After year 9, shortfalls 80-100%	For CONUS, mirrors No Upgrade Case For COCOM, through year 52, shortfalls 15% Years 2-7, shortfalls around 50% After year 7, shortfalls 80-100%
Regional Upgrade (CONUS Upgrades)	For CONUS, mirrors Everyone Upgrades Case For COCOM, through year 8, shortfalls 10-20% Years 9-14, shortfalls around 50% After year 15, exceeded 100%	For CONUS, mirrors Everyone Upgrades Case For COCOM, through year 4 ½, shortfalls 15-30% Years 5- 8, shortfalls 50-80% After year 8, exceeded 100%	For CONUS, mirrors Everyone Upgrades Case For COCOM, through year 3, shortfalls 10-20% Years 4-9, shortfalls 40-80% After year 9, exceeded 100%
Everyone Upgrades Simultaneously	Network exceeded 10-20% until year 5, when demand no longer exceeds capacity	Network exceeded 10-20% until year 5, when demand no longer exceeds capacity	Network exceeded 10-20% until year 5, when demand no longer exceeds capacity